arXiv:1403.7238v1 [cs.DM] 27 Mar 2014

Reduction Techniques for Graph Isomorphism in the Context of Width Parameters

Yota Otachi and Pascal Schweitzer*

Japan Advanced Institute of Science and Technology
School of Information Science
Asahidai 1-1, Nomi, Ishikawa 923-1292, Japan
otachi@jaist.ac.jp

RWTH Aachen University
Ahornstraße 55, 52074 Aachen, Germany
schweitzer@informatik.rwth-aachen.de

March 31, 2014

Abstract

We study the parameterized complexity of the graph isomorphism problem when parameterized by width parameters related to tree decompositions. We apply the following technique to obtain fixed-parameter tractability for such parameters. We first compute an isomorphism invariant set of potential bags for a decomposition and then apply a restricted version of the Weisfeiler-Lehman algorithm to solve isomorphism. With this we show fixed-parameter tractability for several parameters and provide a unified explanation for various isomorphism results concerned with parameters related to tree decompositions.

As a possibly first step towards intractability results for parameterized graph isomorphism we develop an fpt Turing-reduction from strong tree width to the a priori unrelated parameter maximum degree.

1 Introduction

The graph isomorphism problem is the algorithmic task to decide whether two given graphs are isomorphic, i.e., whether there exists a bijection from the vertices of one graph to the vertices of the other graph preserving adjacency and non-adjacency. The problem is situated in the complexity class NP. However, despite extensive research on this problem, the complexity remains unknown. It is neither known whether the problem is polynomial-time solvable nor whether it is NP-hard.

In this paper, we are interested in the parameterized complexity of the isomorphism problem. For other aspects related to the isomorphism problem we refer the reader to other sources (e.g., [23], [28], [31]).

In the parameterized context, for a graph parameter k, such as the maximum degree of the input graphs, we ask for an algorithm that solves isomorphism of graphs with parameter at

*This work is partially supported by Grant-in-Aid for Young Scientists (B) 25730003, Japan Society for the Promotion of Science.

most k. In this context, we are interested in the existence of algorithms with a running time of $O(f(k)n^c)$ for some constant $c \in \mathbb{N}$ in contrast to algorithms with a running time of $O(n^{f(k)})$. Running times of the former type are called fpt time and the algorithms are said to be fixed-parameter tractable algorithms.

Related work. There are various results that show that isomorphism is fixed-parameter tractable with respect to some parameter. Such results exist for the parameters color multiplicity [18] (also known for hypergraphs [1]), eigenvalue multiplicity [13], rooted distance width [34], feedback vertex set number [25], bounded permutation distance [29], tree-depth [7] and connected path distance width [27]. For chordal graphs, tractability results are known for the parameters clique number [22], [26] and the size of simplicial components [32]. Yet, for many parameters, such as maximum degree, tree width and genus, it is not known whether there exist fixed-parameter tractable algorithms solving isomorphism (see [25]). However, no non-tractability results are known. One of the obstacles to understanding the parameterized complexity of graph isomorphism is the uncertainty whether the standard reduction techniques, like showing W[1]-hardness, can be applied (see the discussions in [25] and [34]).

Our results. We study the parameterized complexity of isomorphism with respect to various parameters related to strong tree decompositions. We first develop a method to obtain fixed-parameter tractable algorithms for parameterized graph isomorphism problems. The underlying technique of many results showing such results is to first find a restricted isomorphism invariant family of sets, potential bags, which capture a tree decompositions and to then use these to perform an isomorphism test that uses some form of dynamic programming. It turns out that it is possible to prove that this technique is applicable in general. To prove this general statement, we develop a restricted version of the Weisfeiler-Lehman color refinement algorithm and prove that it successfully decides isomorphism whenever an invariant family of potential bags capturing tree decompositions is available for the input graphs. The algorithm neither computes a decomposition nor is it required that a decomposition is given to the algorithm.

Using the technique, we show tractability of graph isomorphism for the parameters root-connected tree distance width and connected strong tree width. We also provide families of examples showing that neither of the two graph parameters mentioned in the theorem can be bounded by a function of the other. The two tractability results extend results in [17], [27], and [34] also concerned with restricted forms of strong tree decompositions, also answering a question from [34]. Furthermore, with the technique, it is for example also possible to show that graph isomorphism parameterized by the maximum of the length of a longest geodesic cycle and strong tree width or by the maximum of the chordality and degree is fixed-parameter tractable.

In general, our technique provides a unified explanation for the various results [7, 17, 22, 25, 26, 27, 32, 34] all showing that certain restrictions on tree decompositions lead to efficient algorithms for the isomorphism problem. Indeed, all of these approaches can be interpreted as determining some restricted family of potential bags capturing a tree decomposition and then performing some form of dynamic programming to check for isomorphism, that can also be performed by the restricted Weisfeiler-Lehman algorithm. In each of the references above, the dynamic programming is a substantial part of the argumentation, which can now be replaced by the general theorem.

Finally, we show how the technique can be applied to obtain parameterized isomorphism algorithm by exploiting knowledge on the set of potential maximal cliques, of which we already know that it can always be computed in polynomial time in the number of potential maximal cliques.

Our technique also provides a proof of the fact that for graphs of bounded tree width a sufficiently high-dimensional Weisfeiler-Lehman algorithm can be used to determine isomorphism. This fact was first proven by Grohe and Mariño using logic [20] (see also [19]) and provides

to date the fastest running time for isomorphism of bounded tree width graphs. Our proof provides a direct argument for this fact, which does not involve logic. We remark that in his book, Toda [31] also gives a dynamic programming algorithm matching the running time of the algorithm of Grohe and Mariño.

In this paper, we also take a first step towards developing means for some form of intractability result. Specifically, for the isomorphism problem we construct an fpt Turing reduction from strong tree width to the a priori unrelated parameter maximum degree. The existence of this reduction in particular implies that if graph isomorphism is fixed-parameter tractable when parameterized by degree then it is also fixed-parameter tractable when parameterized by strong tree width. However, a possibly better interpretation of this result is that isomorphism parameterized by degree is hard, being at least as intractable as isomorphism parameterized by strong tree width.

To obtain the reduction, we reduce the problem to biconnected components, a technique frequently used for isomorphism algorithms concerned with planar graphs (see [10]). However, we require an extended form of such a reduction allowing us to work with graphs equipped with an equivalence relation and equipped with a coloring of the linear orders of the equivalence classes.

2 Preliminaries

In this paper all graphs are finite, simple, undirected graphs. A *biconnected component* (also called a block) is a maximal connected subgraph not containing a cut-vertex. In particular, the connected graph on 2 vertices is biconnected.

A *strong tree decomposition* of a graph $G = (V, E)$ is a pair $(\{X_i \mid i \in I\}, T = (I, F))$ where $\{X_i \mid i \in I\}$ is a partition of the vertex set V into so-called bags X_i and $T = (I, F)$ is a tree such that the following holds: for all edges $\{u, v\} \in E$, either there is $i \in I$ with $u, v \in X_i$, or there are two adjacent tree vertices $i, i' \in I$ such that $u \in X_i$ and $v \in X_{i'}$. A *connected strong tree decomposition* is a strong tree decomposition for which each bag X_i induces a connected subgraph. The *width* of a strong tree decomposition is the maximum size of a bag of the decomposition.

A strong tree decomposition $(\{X_i \mid i \in I\}, T = (I, F))$ with a distinguished root $r \in I$ is a *tree distance decomposition* if each $v \in X_i$ with $i \neq r$ has a neighbor $u \in X_j$ where j is the parent of i in T rooted at r. A tree distance decomposition with root r is a *root-connected tree distance decomposition* if X_r induces a connected subgraph.

Here, we slightly diverge from the terminology used in [27] to highlight the fact that only the root set must induce a connected graph, and thereby avoid confusion with the term connected strong tree decomposition.

For a class of decompositions $\mathcal{C}$, the $\mathcal{C}$ width of a graph G is the minimal width over all $\mathcal{C}$ decompositions of G. We thus obtain the graph parameters strong tree width, denoted $\mathrm{stw}(G)$, connected strong tree width, denoted $\mathrm{cstw}(G)$, tree distance width, denoted $\mathrm{tdw}(G)$ and root-connected tree distance width, denoted $\mathrm{rctdw}(G)$. The notion of strong tree width was introduced by Seese [30] and is also known as *tree-partition width* [12]. In the context of graph isomorphism, tree distance decompositions were first considered in [34].

For a graph G, there may be several tree distance decompositions with the same root set S. However, there is a unique minimal decomposition (i.e., the partition into bags at least as fine as any other partition into bags obtained from a tree distance decomposition) with root set S. Given S, this minimal decomposition can be computed linear time.

Theorem 1 ([34, Theorem 2.1]). *Given a graph G and a set S, one can compute in $O(m)$ time the unique tree distance decomposition with root set S.*

We denote the width of this decomposition by $\mathrm{tdw}_S(G)$. Note that if G is not connected, it may be the case that there is no tree distance decomposition with root set S. To facilitate our proofs and simplify algorithms, we define $\mathrm{tdw}_S(G)$ to be infinite in this case.

For a graph G with distinct non-adjacent vertices s and t an *s-t-separator* is a set of vertices S such that s and t are in different components of $G - S$. An s-t separator is minimal if no proper subset of S is an s-t-separator.

An *fpt Turing reduction* (see [14]) of a parameterized problem P_1 with parameter k_1 to a parameterized problem P_2 with parameter k_2 is a Turing reduction from P_1 to P_2 with fpt running time for which the parameter k_2 of all oracle calls to the problem P_2 is bounded by a computable function in terms of k_1. In other words, a turing reduction is an fpt-algorithm solving the parameterized problem P_1 with the help of an oracle that solves problem P_2 such that there exists a computable function g such that for all oracle queries $y \in P_2$ posed on an input x with parameter k_1 it holds that the parameter k_2 of y is at most $g(k_1)$.

Suppose we assign to every graph G a subset of the vertices $\mathcal{V}(G) \subseteq V(G)$. We say this assignment is *isomorphism invariant* if for every isomorphism $\pi\colon G_1 \to G_2$ we have $\mathcal{V}(G_2) = \pi(\mathcal{V}(G_1))$. This definition extends to assignments of tuples or sets of vertex sets and also to colored graphs.

3 Tree decompositions and the Weisfeiler-Lehman algorithm

In the graph isomorphism literature, for various graph classes, results are known showing that the Weisfeiler-Lehman algorithm yields polynomial time isomorphism algorithms (see [19]). In this section we describe a restricted version of the Weisfeiler-Lehman algorithm and show that it can be used to obtain fixed-parameter tractability results. Intuitively, the k-dimensional Weisfeiler-Lehman algorithm repeatedly recolors k-tuples of vertices by assigning them a color that depends on the multiset of previous colors of adjacent k-tuples, where tuples are adjacent if they differ by at most one entry. Our restricted version of the algorithm performs this recoloring operation only on a restricted set of k-tuples. For more information on the standard Weisfeiler-Lehman algorithm we refer the reader to existing literature (see [2] and [28] for more pointers).

For $k \geq 2$ we now define the *restricted k-dimensional Weisfeiler-Lehman color refinement*. We say a family of sets $\mathcal{V}$ has *width* k if the largest set in $\mathcal{V}$ has size k. Let G be a graph and $\mathcal{V}$ be a family of sets of vertices of G of width at most k'. Let $\mathcal{V}^+$ be the set of k-tuples $(v_1, \ldots, v_k)$ (with entries not necessarily distinct) for which $\{v_1, \ldots, v_{k'}\}$ is in $\mathcal{V}$. For every k-tuple $(v_1, \ldots, v_k)$ in $\mathcal{V}^+$ we define $\mathrm{wl}_0^k[\mathcal{V}, G](v_1, \ldots, v_k)$ as the isomorphism type of the subgraph induced by the ordered tuple $(v_1, \ldots, v_k)$. If the graph is colored then the isomorphism type has to take the coloring into account. More precisely, the coloring wl_0^k is a coloring that satisfies $\mathrm{wl}_0^k[\mathcal{V}, G](v_1, \ldots, v_k) = \mathrm{wl}_0^k[\mathcal{V}, G](v'_1, \ldots, v'_k)$ if and only if we can map v_i to v'_i and obtain an isomorphism of the colored graphs induced by $\{v_1, \ldots, v_k\}$ and $\{v'_1, \ldots, v'_k\}$. If $(v_1, \ldots, v_k) \notin \mathcal{V}^+$ then we define $\mathrm{wl}_0^k[\mathcal{V}, G](v_1, \ldots, v_k)$ to be the empty set $\emptyset$.

Iteratively for $i \geq 0$, we define $\mathrm{wl}_{i+1}^k[\mathcal{V}, G](v_1, \ldots, v_k)$ to be the empty set $\emptyset$ if $(v_1, \ldots, v_{k'}) \notin \mathcal{V}^+$ and to be $\big(\mathrm{wl}_i^k[\mathcal{V}, G](v_1, \ldots, v_k), \mathcal{M}_i^k\big)$ otherwise, where $\mathcal{M}_i^k$ is the multiset given by

$$\mathcal{M}_i^k := \{\!\{(\mathrm{wl}_i^k[\mathcal{V}, G](x, v_2, \ldots, v_k), \mathrm{wl}_i^k[\mathcal{V}, G](v_1, x, v_3, \ldots, v_k), \ldots,$$

$$\mathrm{wl}_i^k[\mathcal{V}, G](v_1, \ldots, v_{k-1}, x)) \mid x \in V(G)\}\!\}.$$

The process partitions the ordered k-tuples into classes according to their color. Since in each iteration the color of the previous iteration is encoded in the new color, k-tuples which are assigned different colors will continue to have different colors in all subsequent iterations. Therefore the refinement process stabilizes. We define $\mathrm{wl}^k_\infty[\mathcal{V}, G](v_1, v_2, \ldots, v_k)$ as $\mathrm{wl}^k_i[\mathcal{V}, G](v_1, v_2, \ldots, v_k)$ where i is the least positive integer such that the induced partition in step i is equivalent to the induced partition in step $i+1$. Abusing notation, we may drop the specifications $[\mathcal{V}, G]$ whenever they are apparent from the context.

Lemma 2. *Suppose for $i \in \{1, 2\}$ we are given a graph G_i and a family of subsets of the vertices $\mathcal{V}_i$ of width $k' \leq k$. For the restricted k-dimensional Weisfeiler-Lehman color refinement the following properties hold:*

1. *If $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k) \neq \emptyset$ holds then also for all $j_1, j_2 \in \{1, \ldots, k\}$ it holds that*

$$\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_{j_1-1}, v_{j_2}, v_{j_1+1}, \ldots, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_{j_1-1}, v'_{j_2}, v'_{j_1+1}, \ldots, v'_k).$$

2. *If $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k) \neq \emptyset$ then for all indices $j_1, j_2 > k'$, if v_{j_1} and v_{j_2} are contained in the same connected component of $G_1 - \{v_1, \ldots, v_{k'}\}$ then v'_{j_1} and v'_{j_2} are contained in the same connected component of $G_2 - \{v'_1, \ldots, v'_{k'}\}$.*

Proof. (Part 1.) If $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k) \neq \emptyset$ then by definition of the stable refinement, there exists a vertex $x \in V_2$ such that the equation

$$\begin{aligned}&\big(\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_{j_2}, v_2, \ldots, v_k), \mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, v_{j_2}, v_3, \ldots, v_k), \ldots, \mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_{k-1}, v_{j_2})\big) = \\ &\big(\mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](\ x\ , v'_2, \ldots, v'_k), \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ x\ , v'_3, \ldots, v'_k), \ldots, \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_{k-1}, \ x\ )\big)\end{aligned}$$

holds. Since wl^k_0 in particular encodes the isomorphism type of the graph induced by its entries, we conclude that $x = v'_{j_2}$, which proves the first part of the lemma.

(Part 2.) We show the following statement by induction on t. If there is a path from v_{j_1} to v_{j_2} in $G_1 - \{v_1, \ldots, v_{k'}\}$ of length at most t, but there is no path of length at most t from v'_{j_1} to v'_{j_2} in $G_2 - \{v'_1, \ldots, v'_{k'}\}$, then $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) \neq \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k)$. For $t = 0$, if there is said path of length 0 from v_{j_1} to v_{j_2} this implies that $v_{j_1} = v_{j_2}$ and that $v_{j_1} \notin \{v_1, \ldots, v_{k'}\}$.

If there is no path of length 0 from v'_{j_1} to v'_{j_2} in $G_2 - \{v'_1, \ldots, v'_{k'}\}$ then $v'_{j_1} \neq v'_{j_2}$ or $v'_{j_1} \in \{v'_1, \ldots, v'_{k'}\}$. Either way, we have that $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) \neq \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k)$ since the initial coloring wl^k_0 encodes the isomorphism type of its entries.

Suppose now the statement has been shown for length $t' < t$ and suppose a shortest path from v_{j_1} to v_{j_2} in $G_1 - \{v_1, \ldots, v_{k'}\}$ is of length $t > 0$. If $v'_{j_1} \in \{v'_1, \ldots, v'_{k'}\}$ or $v'_{j_1} \in \{v'_1, \ldots, v'_{k'}\}$ the statement follows as in the induction base since wl^k_0 encodes the isomorphism type of its entries. Also note that $t > 0$ implies $v_{j_1} \neq v_{j_2}$. Let x be a vertex in $G_1 - \{v_1, \ldots, v_{k'}\}$ which is adjacent to v_{j_2} and of distance $t - 1$ from v_{j_1}. For every vertex y, either v'_{j_2} is not adjacent to y or there is no path of length at most $t - 1$ from v'_{j_1} to y in $G_2 - \{v'_1, \ldots, v'_{k'}\}$. Thus, by induction, we conclude that it cannot be simultaneously the case that $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_{j_1-1}, x, v_{j_1+1}, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_{j_1-1}, y, v'_{j_1+1}, v'_k)$ and also that $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_{j_2-1}, x, v_{j_2+1}, v_k) = \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_{j_2-1}, y, v'_{j_2+1}, v'_k)$. This shows that $\mathrm{wl}^k_\infty[\mathcal{V}_1, G_1](v_1, \ldots, v_k) \neq \mathrm{wl}^k_\infty[\mathcal{V}_2, G_2](v'_1, \ldots, v'_k)$. □

Lemma 3. *For a graph G and a family $\mathcal{V}$ sets of vertices of G of width k', the stable partition of the restricted $(k'+c)$-dimensional Weisfeiler-Lehman color refinement can be computed in time $\mathcal{O}\big((k'+c)^2 \cdot |\mathcal{V}^+| n \cdot \log(|\mathcal{V}^+|)\big)$.*

Proof. The technique in this proof, computing the stable partition in said running time, is an adaptation of the technique of Immerman and Lander [21] (see also [2]) that computes the stable refinement of the (unrestricted) Weisfeiler-Lehman algorithm in time $\mathcal{O}(k^2 n^{k+1} \log(n))$. We define $k = k' + c$. If the given graph is uncolored, the initial coloring of $\mathrm{wl}_0^k(v_1, \ldots, v_k)$ can be chosen to be the string $E(v_1, v_1), \ldots, E(v_1, v_k), E(v_2, v_1), \ldots, E(v_k, v_k)$, where $E(v_i, v_j) = 1$ if (v_i, v_j) is an edge and 0 otherwise. If the graph is already equipped with a coloring, we define an initial coloring as the pair of the color just defined and the initially given color. Given the initial coloring we proceed as follows: We again let $\mathcal{V}^+$ be the set of all k-tuples $(v_1, \ldots, v_k)$ for which $\{v_1, \ldots, v_{k'}\}$ is in $\mathcal{V}$. We use a list S that will contain subsets of $\mathcal{V}^+$. Initially it contains the color classes according to the initial coloring. The algorithm now repeatedly splits color classes using the first element of the list, which is removed afterwards. More specifically, suppose B is the first element of the list. Let B' be the elements of $\mathcal{V}^+$ which can be obtained by replacing up to one vertex in an element of B. For each element b' of B' we compute a new color using a modified Weisfeiler-Lehman recursion. That is, if χ is the previous coloring, the new color of $b' = (b_1, \ldots, b_k)$ will be $\chi'(b') = \big(\chi(b'), \mathcal{M}_i^k\big)$, where $\mathcal{M}_i^k$ is the multiset given by

$$\{\!\{(\chi(x, b_2, \ldots, b_k), \ldots, \chi(b_1, \ldots, b_{k-1}, x)) \mid x \in V(G) \text{ and there}$$

$$\text{exists } i \in \{1, \ldots, k\} \text{ such that } (b_1, \ldots, b_{i-1}, x, b_{i+1}, \ldots, b_k) \in B\}\!\}.$$

For each color class that is split into several smaller color classes during the procedure, we add all but the largest of these smaller color classes to the end of the list S. The algorithm continues until the list is empty. The stable partition is the partition induced by the final color classes. To prevent the names of colors from becoming excessively long strings, we always rename newly arising colors by assigning previously unused integers. Details are given in Algorithm 1.

Correctness. The computed partition is coarser than or equal to the stable partition of the restricted Weisfeiler-Lehman algorithm, since each step splits color classes according to color induced subsets of the multisets from in the definition of the restricted Weisfeiler-Lehman algorithm using previously computed colors. For these previously computed colors, the partition is coarser than or equal to the stable partition by induction. For correctness it thus suffices to argue that when given the final coloring, which we denote by χ, the restricted Weisfeiler-Lehman algorithm does not refine the color classes.

Supposing otherwise, there must be tuples $(b_1, \ldots, b_k)$ and $(b'_1, \ldots, b'_k)$ in $\mathcal{V}^+$, such that the restricted Weisfeiler-Lehman algorithm assigns these tuples different colors in the first iteration when provided with χ as initial colors. Thus, there must be a k-tuple of color classes $(B_1, \ldots, B_k)$ such that the sets $\{x \in V(G) \mid \chi(x, b_2, \ldots, b_k) \in B_1, \ldots, \chi(b_1, \ldots, b_{k-1}, x) \in B_k\}$ and $\{x \in V(G) \mid \chi(x, b'_2, \ldots, b'_k) \in B_1, \ldots, \chi(b'_1, \ldots, b'_{k-1}, x) \in B_k\}$ have different cardinality. Let B_i be the color class among $\{B_1, \ldots, B_k\}$ for which at the latest time during the algorithm there is some superset $\hat{B}$ in the list S. If this set is B_i then Algorithm 1 would have distinguished $(b_1, \ldots, b_k)$ and $(b'_1, \ldots, b'_k)$. This shows that B_i was the largest class among the classes into which the superclass $\hat{B}$ was split. However, in this case, among the classes into which $\hat{B}$ was split, there is a class that can replace B_k in the sequence $(B_1, \ldots, B_k)$ yielding a new sequence for which we also obtain two sets of different cardinality.

Running time. To bound the running time, it suffices to bound the amount of work performed in the for-loop of Algorithm 1. Note that whenever a set is added to the list, its size is at most half of the size of the class that is split. Thus every element can be in at most $\mathcal{O}(\log(|\mathcal{V}^+|))$ many of the sets which are at some point in S. It suffices now to show that it is possible to compute $\chi'(b')$ for all elements in B' in time $\mathcal{O}((k' + c)^2 \cdot |B'|n)$. Note that every element in the multiset used to compute $\chi'(b')$ contains at least one coordinate that is in B. We can therefore compute $\chi'(b')$ for all elements in B' by determining the counts of elements in said multiset as

Algorithm 1 The Weisfeiler-Lehman algorithm restricted to $\mathcal{V}$

Input: Integers k' and c, a graph G, a family $\mathcal{V}$ of sets of vertices of G of width k', and a possibly uniform coloring χ of all k-tuples in $\mathcal{V}^+$ (where $k = (k' + c)$).

Output: The stable partition of the restricted $(k' + c)$-dimensional Weisfeiler-Lehman color refinement.

1: for all tuples $(v_1, \ldots, v_k) \in \mathcal{V}^+$ define the color $\mathrm{wl}_0^k(v_1, \ldots, v_k)$ as $(\chi(v_1, \ldots, v_k), (E(v_1, v_k), E(v_2, v_1), \ldots, E(v_k, v_k)))$; for all other tuples the initial coloring is $\emptyset$
2: initialize an empty list S
3: add each color class in $\mathcal{V}^+$ to the list in lexicographic order
4: **while** S is not empty **do**
5: let B be the first set of S
6: let B' be the set of elements $b' \in \mathcal{V}^+$ for which there exists a set $b \in B$ differing from b' by at most one vertex
7: compute $\chi'(b')$ for all $b' \in B'$
8: replace χ with χ' for elements in B'
9: for all color classes that are split due to this replacement of χ add all but the largest of the new color classes to the end of the list S
10: remove B from the front of the list
11: **end while**
12: **return** the partition induced by χ

follows. We iterate over all elements b of B and all coordinate positions j in $\{1, \ldots, k\}$ and add suitable counts to all elements b' in B' for which we can obtain b by replacing coordinate j of b' with coordinate j of b. The k-tuple of colors which is to be counted is obtained by inserting this value of coordinate j of tuple b one by one into the other coordinates of b'. To avoid double counting, the corresponding counter is then incremented if j is the least position for which the entry is in B. Iterating over the elements in each color class that is split, we can in each case find the largest color class. In time linear in the sizes of the split color classes. □

We remark that when the described efficient version of the restricted Weisfeiler-Lehman algorithm is used on two graphs for comparing colors of tuples between the two graphs, some care has to be taken concerning hashing. If hashing of colors is used, then the algorithm has to be performed on both graphs at the same time (or alternatively on the disjoint union of the two graphs) to avoid hashing different color to the same value.

Being a restriction implies that the known examples that cannot be solved by the Weisfeiler-Lehman algorithm [8] can also not be solved by the restricted version. However, we will now prove that the restricted Weisfeiler-Lehman algorithm decides isomorphism of graphs whenever the set $\mathcal{V}(G)$ captures a tree decomposition. To facilitate the proof and to make it more easily applicable in the future, we prove the theorem for tree decompositions, instead of proving it just for strong tree decompositions.

Recall that a *tree decomposition* is a pair $(\{X_i \mid i \in I\}, T = (I, F))$ for which $\bigcup_{i \in I} X_i = V(G)$ and $T = (I, F)$ is a tree such that every vertex is contained in some bag, for adjacent vertices there is a bag containing both of them and for every vertex v the set of bags containing v induces a connected subtree of T.

Given a graph G we say that a family of sets $\mathcal{V}(G)$ *captures a tree decomposition* $\mathcal{T}$ of G if every bag is in $\mathcal{V}(G)$. If G is equipped with an equivalence relation and possibly a tuple-coloring, we additionally require that every equivalence class is contained in a bag of $\mathcal{T}$. We say that

a tree decomposition is *semi-smooth* if the intersection of adjacent bags has size at most one smaller that the size of the larger bag (for the decomposition to be smooth one also requires that all bags have the same size, see [3]).

Theorem 4. *Suppose we are given an algorithm that computes for every graph G in a graph class $\mathcal{C}$ an isomorphism invariant family of vertex sets $\mathcal{V}(G)$ of width at most k' such that $\mathcal{V}(G)$ captures a semi-smooth tree decomposition of G. Then we can decide isomorphism of graphs in $\mathcal{C}$ with the $(k'+3)$-dimensional Weisfeiler-Lehman algorithm restricted to $\mathcal{V}(G)$.*

Proof. Suppose G and G' are graphs in $\mathcal{C}$. We run the Weisfeiler-Lehman algorithm and claim the graphs are isomorphic if and only if the multisets of colors obtained as the stable refinement of each graph are isomorphic.

Since functor $\mathcal{V}(\cdot)$ is isomorphism invariant, the restricted Weisfeiler-Lehman algorithm is also isomorphism invariant. Therefore, if G and G' are isomorphic then, this procedures claims that the graphs are isomorphic.

We now show that if the restricted Weisfeiler-Lehman algorithm terminates with the multisets of colored tuples being equal in both graphs, then the graphs are isomorphic.

Let $\mathcal{T}$ be a semi-smooth tree decomposition of G that is captured by $\mathcal{V}(G)$. Let $k = k' + 3$. We now prove the following statement by induction on t:

Claim 5. *If $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, z, z) = \mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', z', z')$ and if the set $\{v_1, \ldots, v_{k'}\}$ is a bag of $\mathcal{T}$, then there is an isomorphism ϕ from the subgraph of G induced by the vertices of $\{v_1, \ldots, v_{k'}\}$ and all vertices in components of $G \setminus \{v_1, \ldots, v_{k'}\}$ not containing z of size at most t to the subgraph of G' induced by the vertices of $\{v'_1, \ldots, v'_{k'}\}$ and all vertices in components of $G' \setminus \{v'_1, \ldots, v'_{k'}\}$ not containing z' of size at most t. Moreover the isomorphism ϕ can be chosen such that v_i maps to v'_i for $i \in \{1, \ldots, k\}$ and such that it respects colors of equivalence classes.*

Under the assumption, the multiset $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, z, z)$ is non-empty, which implies that the set $\{v'_1, v'_2, \ldots, v'_{k'}\}$ is contained in an element of $\mathcal{V}(G')$.

For $t = 0$, by definition $\mathrm{wl}^k_0(v_1, v_2, \ldots, v_{k'}, z, z, z) = \mathrm{wl}^k_0(v'_1, v'_2, \ldots, v'_{k'}, z', z', z')$ if and only if the map sending v_i to v'_i and z to z' is an isomorphism from the graph $G[\{v_1, \ldots, v_k, z\}]$ to the graph $G'[\{v'_1, \ldots, v'_k, z'\}]$. This resolves the base case.

Suppose now $t > 0$ and that $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, z, z) = \mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', z', z')$ where the set $\{v_1, \ldots, v_{k'}\}$ is a bag of $\mathcal{T}$.

Let $B = \{v_1, \ldots, v_{k'}\}$and $B' = \{v'_1, \ldots, v'_{k'}\}$. We call the components of $G - B$ and $G' - B'$ that do not contain z or z', respectively, and which are of size at most t the *components of interest*. Let C be a component of interest in $G - B$. Since $\mathcal{T}$ is semi-smooth, there is a bag $\hat{B}$ in $\mathcal{T}$ containing a vertex w of C and all but one of the vertices from B. Moreover there is a vertex in $\{v_1, \ldots, v_{k'}\}$ such that no vertex of C different from w is adjacent to this vertex. Without loss of generality we assume that v_1 fulfills this property.

Since $\mathrm{wl}^k_\infty(v_1, \ldots, v_{k'}, z, z, z) = \mathrm{wl}^k_\infty(v'_1, \ldots, v'_{k'}, z', z', z')$ there must be a vertex w' in G' such that $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, w, z) = \mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', w', z')$ holds. This implies that $w' \notin \{v'_1, \ldots, v'_{k'}\}$. By Lemma 2, this also implies that w' is not in the same component as z' in $G' - \{v'_1, \ldots, v'_{k'}\}$. Let C' be the component of $G' - \{v'_1, \ldots, v'_{k'}\}$ containing w'. We show that the map that sends v_i to v'_i can be extended to an isomorphism that maps C to C'. Since by Lemma 2 the equation $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, w, z) = \mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', w', z')$ in particular implies $\mathrm{wl}^k_\infty(w, v_2, \ldots, v_{k'}, v_1, v_1, v_1) = \mathrm{wl}^k_\infty(w', v'_2, \ldots, v'_{k'}, v'_1, v'_1, v'_1)$, the map sending w to w' and v_i to v'_i for $i \geq 2$ extends to an isomorphism mapping all components of $G - \{w, v_2 \ldots, v_{k'}\}$ of size at most $t - 1$ not containing v_1 to components of $G - \{w', v'_2 \ldots, v'_{k'}\}$ not containing v'_1. Since C and C' are connected and contain w and w' respectively, this implies that this mapping maps C to C'.

We summarize what we have achieved so far towards proving the claim: for every component C of $G - B$ of interest (i.e., of size at most t and not containing z) we have found a component C' of $G'-B'$ of interest (i.e., of size at most t and not containing z') such that $G[C\cup B]$ can map isomorphically to $G'[C \cup B']$. We now need to show that we can find an isomorphism that simultaneously maps all components of interest in $G - B$ bijectively to the components of interest in $G' - B'$.

From what we have shown so far, it could in principle be possible that by choosing a different vertex of C as w, the new corresponding vertex w' would have to be taken outside of C' and thus give us a different component C'' of $G' - B'$. To argue that this is not the case, we show that $\{\{\mathrm{wl}^k_\infty(v_1, \ldots, v_{k'}, z, u, z) \mid u \in C\}\} = \{\{\mathrm{wl}^k_\infty(v'_1, \ldots, v'_{k'}, z', u', z') \mid u' \in C'\}\}$ as multisets. To show this, it suffices by Lemma 2 to show that $\{\{\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, u, w) \mid u \in C\}\} = \{\{\mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', u', w') \mid u' \in C'\}\}$.

Since $\mathrm{wl}^k_\infty(v_1, \ldots, v_{k'}, z, w, z) = \mathrm{wl}^k_\infty(v'_1, \ldots, v'_{k'}, z', w', z')$ it follows also from Lemma 2 that $\mathrm{wl}^k_\infty(w, v_2, \ldots, v_{k'}, z, z, w) = \mathrm{wl}^k_\infty(w', v'_2, \ldots, v'_{k'}, z', z', w')$ and consequently it follows that as multisets $\{\{\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, u, w) \mid u \in V\}\} = \{\{\mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', u', w') \mid u' \in V\}\}$. It suffices now for us to show that if ($x \in C$ and $x' \notin C'$) or if ($x \notin C$ and $x' \in C'$) then automatically also the inequality $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_{k'}, z, x, w) \neq \mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_{k'}, z', x', w')$ holds. This however follows again by Lemma 2.

We define the sought isomorphism. Since $\mathrm{wl}^k_\infty(v_1, \ldots, v_{k'}, z, z, z) = \mathrm{wl}^k_\infty(v'_1, \ldots, v'_{k'}, z', z', z')$, we can find isomorphism from the components of $G-B$ of interest to the components of $G'-B'$ of interest such these isomorphism are compatible with the map from B to B' mapping v_i to v'_i. This yields an isomorphism from the graph induced by all vertices in B and all vertices in components of interest bijectively to those of G'.

To finish the proof of the claim, it remains to argue that the isomorphism we constructed respects the coloring of the graphs. Let $(v_1, \ldots, v_t)$ be an ordering of an equivalence class of the vertices of G. Since there is a bag of $\mathcal{V}(G)$ that contains the equivalence class,the multiset $\mathrm{wl}^k_\infty(v_1, v_2, \ldots, v_t, v_t, \ldots, v_t)$ is not empty. Moreover under the isomorphism we have constructed, the color of the image, $\mathrm{wl}^k_\infty(v'_1, v'_2, \ldots, v'_t, v'_t, \ldots, v'_t)$ must be non-empty and thus the induced subgraph $G[v_1, \ldots, v_t]$ must map to the induced subgraph $G'[v'_1, \ldots, v'_t]$ respecting the coloring of the equivalence class. By symmetry, colored orderings of equivalence classes of G' also map to orderings of equivalence classes of G of the same color. This proves the claim.

We now apply the claim to finish our proof. By setting z equal to $v_{k'}$ and t equal to $n + 1$ we obtain from the claim the existence of an isomorphism that maps all vertices of G bijectively to the vertices of G'. □

The previous theorem requires $\mathcal{V}(G)$ to capture a semi-smooth tree decomposition. However, we can extend the theorem to tree decompositions and strong tree decompositions by using the alternative set $\mathcal{V}' = \{B_1 \cup B_2 \mid B_1, B_2 \in \mathcal{V}(G)\}$. This can be seen by the following two observations.

If $\mathcal{V}$ is isomorphism invariant and captures a tree decomposition of G, then $\mathcal{V}'$ is isomorphism invariant and captures a semi-smooth tree decomposition. This follows from the construction that produces a smooth tree decomposition from a tree decomposition given in [3].

Suppose $\mathcal{B}$ is the set of bags of a strong tree decomposition. Then there is a semi-smooth tree decomposition $\mathcal{B}'$ such that for every bag B of $\mathcal{B}'$ there are bags B_1 and B_2 in $\mathcal{B}$ such that $B \subseteq B_1 \cup B_2$. This can be seen with the standard way of constructing a tree decomposition from a strong tree decomposition by inserting for each edge between two bags B_1 and B_2 a path of bags transforming B_1 to B_2 by replacing successively one vertex after the other. We conclude, if $\mathcal{V}$ captures a strong tree decomposition then $\mathcal{V}'$ captures a semi-smooth tree decomposition. This also shows that by setting $\mathcal{V}$ to be the set of all k-tuples of vertices, every graph of tree width

at most k has a smooth tree decomposition captured by $\mathcal{V}'$ and shows that the a sufficiently high-dimensional Weisfeiler-Lehman algorithm solves graph isomorphism of graphs of bounded tree width, as mentioned in the introduction.

Concerning the previous two observations, in our context relevant is the following fact. If we start with a family $\mathcal{V}$ of fpt size in some parameter k with the maximum set size bounded by a function of k and we apply either construction, we obtain a family $\mathcal{V}'$ that is also of fpt size.

Corollary 6. *For a parameter k', given an fpt-algorithm that computes for every graph G in a graph class $\mathcal{C}$ an isomorphism invariant family of vertex sets $\mathcal{V}(G)$ of width at most k' such that $\mathcal{V}(G)$ captures a tree decomposition (or a strong tree decomposition), isomorphism of graphs in $\mathcal{C}$ is fixed parameter tractable in k'.*

Proof. Since the algorithm computing $\mathcal{V}$ is an fpt-algorithm, the size of $\mathcal{V}$ is bounded by $f(k')n^d$ for some function f and constant $d \in \mathbb{N}$. The set $\mathcal{V}'$ as defined above has thus size at most $f(k')^2 n^{2d}$, width at most $2k'$, and captures a smooth tree decomposition. Theorem 4 shows that it suffices to perform the $(2k'+3)$-dimensional Weisfeiler-Lehman algorithm restricted to $\mathcal{V}'$ to solve isomorphism, which can be done im fpt time by Lemma 3. □

We remark that if we were interested in actual running times, in the tree width case it is possible to avoid that the increase in width from $\mathcal{V}$ to $\mathcal{V}'$ yielding better running times bounds.

4 Tree distance decompositions with connected root bags

As a first application we show that graph isomorphism parameterized by root connected tree distance width is fixed parameter tractable. We say two vertices v_1 and v_2 in a graph are *k-connected*, if there are k internally vertex disjoint paths from v_1 to v_2. We denote this by $v_1 \equiv_k v_2$. The task of checking whether two vertices are k-connected is also known as the Menger Problem and can be solved in polynomial time via a reduction to the maximum flow problem.

Lemma 7. *Let G be a graph containing vertices v_1 and v_2. If $v_1 \equiv_{2k} v_2$, then in every strong tree decomposition of width at most k the vertices v_1 and v_2 are in the same bag.*

Proof. Suppose v_1 and v_2 are in different bags B_v and $B_{v'}$ of a strong tree decomposition of G of width at most k. Let B' be the bag adjacent to B_v in the tree along the path from B_v to $B_{v'}$. (It is possible that $B_v = B'$.) Suppose there are $2k$ internally vertex-disjoint paths from v_1 to v_2. Each of these paths must contain an edge with an endpoint in B_v and an endpoint in B'. Let M be a set of $2k$ such edges each belonging to a different path. For every two edges in M, the end vertices must be distinct unless they are v_1 or v_2. Since the bag B_v has size at most k there are at most k edges in M that have v_1 as an endpoint. Moreover there are at most $k-1$ edges in M that have an endpoint in B_v different from v_1. This shows that M has size at most $2k-1$, contradicting the existence of $2k$ internally vertex-disjoint paths between v_1 and v_2. □

A similar result for tree decompositions, stipulating the existence of a bag that contains both v_1 and v_2 in tree decompositions of width at most k whenever $v_1 \equiv_{k+1} v_2$, can be found in [4].

Suppose we are searching for a root set S and have already found a subset S' of S. In this case, the previous lemma provides means of finding other vertices that must also be contained in S. The following lemmas are also concerned with extending S' to a suitable root set.

Lemma 8. *Let G be a connected graph and let S be a root set such that $\mathrm{tdw}_S(G) = k$. Let S' be a subset of S. Let C be a component of $G - S'$.*

1. *If* $\mathrm{tdw}_{S'}(G[S' \cup C]) > k$ *then* $(S - S') \cap C$ *is non-empty.*

2. *If* $\mathrm{tdw}_{S'}(G[S' \cup C]) \leq k$ *then* $\mathrm{tdw}_{S \setminus C}(G) \leq k$.

Proof. For the first part, if $(S - S') \cap C$ were empty then the every tree distance decomposition on G with root set S induces a tree distance decomposition on $G[S' \cup C']$ with root set S', which contradicts $\mathrm{tdw}_{S'}(G[S' \cup C]) > k$.

For the second part let $\mathcal{S}$ be a tree distance decomposition of G with root set S of width at most k. Let $\mathcal{S}'$ be a tree distance decomposition of $G[S' \cup C]$ with root set S' of width at most k. We construct a new tree distance decomposition of G: The set of bags is the union of the sets $M_1 = \{B \setminus C \mid B \in \mathcal{S}\}$ and $M_2 = \mathcal{S}' \setminus \{S'\}$. The root bag of the decomposition is $S \setminus C$. It suffices now to show that this is a rooted tree distance decomposition of width at most k. Note that, together, the sets M_1 and M_2 partition the vertices of G. To see that the graph induced by this partitioning is a tree, first note that all neighbors of a vertex contained in a set in M_2 which are contained in a bag in M_1 are vertices in S'. Since $S \setminus C \supset S'$ is a bag, it suffices thus to show that M_1 and $M_2 \cup \{S'\}$ form strong tree decompositions on $G \setminus C$ and $G[C \cup S']$ respectively. For the latter this follows since it constitutes the original decomposition of $G[S' \cup C]$. For the former this follows since the graph is connected and the new bags are subsets of distinct former bags, and therefore the formed decomposition cannot have cycles. Since all bags of the new decomposition are subsets of previous bags the new decomposition has width at most k. □

Lemma 9. *Let G be a graph and S be a root set such that $\mathrm{tdw}_S(G) = k$. Let S' be a subset of S. If C is a component of $G - S'$ with $|N(S') \cap C| > \ell$ for some positive integer $\ell \in \mathbb{N}$ then there exists $v \in C$ which is connected by more than $(\ell - k)/k^2$ internally vertex-disjoint paths to S'.*

Proof. If $\ell \leq k$ there is nothing to show, so we assume otherwise. Consider a strong tree decomposition of width at most k with root set S and consider the components $C_1, \ldots, C_t$ of $G - S$ that contain a vertex of $N(S') \cap C$. Let C_i be such a component. Since every bag contains at most k vertices, and since vertices in $N(S) \cap C_i$ must be in the same bag, each component C_i contains at most k vertices of $N(S') \cap C$. Since $|N(S') \cap C| > k$, this implies that there must be some vertex in $v \in S \setminus S'$ that has a neighbor in C_i. In particular, we can find a path from some vertex in S' to v whose internal vertices all lie in C_i. Each vertex in $N(S') \cap C$ must either be contained in S or in a component C_i for some $i \in \{1, \ldots, t\}$. Since each C_i can only contain k vertices of $N(S') \cap C$ and since $|N(S') \cap C| > \ell$ there must be more than $(\ell - k)/k$ components. This leads to more than $(\ell - k)/k$ internally vertex-disjoint paths from vertices in S' to vertices in $S \setminus S'$. By the pigeonhole principle there is a vertex $v \in S \setminus S'$ connected by more than $((\ell - k)/k)/k = (\ell - k)/k^2$ internally vertex-disjoint paths to S'. Note that, since C is a connected component, the vertex v is in C. □

The three preceding lemmas can be assembled to find suitable root sets, if they exist, allowing for efficient isomorphism tests.

Theorem 10. *The isomorphism problem parameterized by root-connected tree distance width can be solved in fpt time.*

Proof. As explained in the introduction given a set S we can compute in polynomial time a unique minimal distance decomposition with root set S ([34, Theorem 2.1]). By simply inspecting the size of all bags, we can also check in polynomial time whether this decomposition has width at most k.

By Corollary 6 it thus suffices for us to show that we can enumerate in fpt time an isomorphism invariant family of root sets that contains at least one root set that yields a tree distance

Algorithm 2 An fpt algorithm for isomorphism parameterized by root-connected tree distance width $\mathrm{ISO}(G, S')$

Input: A graph G a subset of the vertices S' and a parameter k.

Output: A non-empty isomorphism invariant list of sets of S for which $S \supset S'$ for which $\mathrm{tdw}_S(G) \leq k$, or report **false**, if no such set exists.

```
1: if S' = {} then
2:   for all v ∈ V(G) do
3:     S' ← v
4:     ISO(G, S')
5:   end for
6: else
7:   while there is a vertex v ∉ S' from which there are 2k internally vertex disjoint paths
     to S' do
8:     S' ← S' ∪ {v}.
9:   end while
10:  if tdw_S(G) ≤ k then
11:    print S
12:  else
13:    if |S'| < k then
14:      let C_1, ..., C_t be the components C of G \ S' for which tdw_{S'}(G[S' ∪ C]) > k
15:      if t + |S'| ≤ k and |N(S') ∩ C_i| ≤ 2k^3 + k for all i ∈ {1, ..., t} then
16:        for all (v_1, ..., v_t) ∈ N(S') ∩ C_1 × ··· × N(S') ∩ C_t do
17:          S' ← S' ∪ {v_1, ..., v_t}
18:          ISO(G, S')
19:        end for
20:      end if
21:    end if
22:  end if
23: end if
```

width of at most k. Indeed for each such root set we can compute the minimal distance decomposition and collect all possible bags of decompositions of width at most k obtained this way. The family of sets obtained has size at most $f(k)n^d$ for some function f and some integer d and captures a strong tree decomposition.

We now describe an algorithm for the task of enumerating the family of root sets, details are given in Algorithm 2. We remark however, that the suitable root set computed by Algorithm 2 is not necessarily connected and the parameter root-connected tree distance width is used only in the analysis of the running time. To enumerate root sets, we start with an arbitrary vertex which we add to a previously empty set S'. While there is a vertex v' not in S' from which there are $2k$ internally vertex disjoint paths to S' we add v' to S'.

If S' induces a tree distance decomposition of width at most k we output S'. Otherwise, let $C_1, \ldots, C_t$ be the components C of $G \setminus S'$ for which $\mathrm{tdw}_{S'}(G[S' \cup C]) > k$. If $t + |S'| > k$ or $|N(S') \cap C_i| > 2k^3 + k$ for some $i \in \{1, \ldots, t\}$ we assess S' as infeasible and backtrack. In case we do not asses S' as infeasible, we add from each component C_i one every vertex contained in $N(S') \cap C_i$ to S' and recurse. By branching, we try all possibilities of adding vertices to S' in this way. That is, for each possible combination of picking one vertex from each $N(S') \cap C_i$ we recursively continue the process.

Whenever this process yields a set S' that induces a tree distance decomposition of width at

most k we declare S' as a candidate and output it.

Since the vertices that are added to S' are always added from a set of vertices defined by an isomorphism invariant property, the sets of candidates computed by the algorithm is isomorphism invariant.

Correctness. To prove the correctness of the algorithm, it remains to show that this algorithm produces a root set that yields a tree distance width of at most k. Let S be a connected set with $\mathrm{tdw}_S(G) \leq k$. Let S'' be a maximal subset of S among the sets S' generated by the algorithm. We show that $\mathrm{tdw}_{S'}(G) \leq k$. For contradiction, we assume otherwise. Since the algorithm initially starts with an arbitrary vertex, S'' is non-empty. Consider the quotient graph in which the vertices of S' have been identified. By Lemma 7, if there is a vertex v not in S'' from which there are $2k$ internally vertex-disjoint paths to S'' then v must be in S. The algorithm will add v_2 to S''. Moreover all vertices added to S'' due their connectivity must be contained in S. Thus the algorithm produces a larger subset contained in S contradicting maximality of S''. If on the other hand no such vertex v_2 exists, by Lemma 8 there exists at least one component (Part 2), but at most k (Part 1) components C with $\mathrm{tdw}_{S''}(G[S'' \cup C]) > k$. By Lemma 9 for each such component $|N(S'') \cap C| \leq 2k^3 + k$. Since S is connected, for each such component there is a vertex in $N(S'') \cap C \cap S$. Thus there is a choice for the algorithm of one vertex for each component that yields a larger subset of S contradicting maximality of S''.

Running time. The running time of Algorithm 2 is $\mathcal{O}(k^{2k^3+k}n^c)$ since each call of the algorithm can be performed in polynomial time, and every vertex, except the first, which is to be added to S' is chosen from a set of size at most $2k^3 + k$ bounding the number of recursive calls to the algorithm by $\mathcal{O}(k^{2k^3+k}n)$. □

We remark that instead of using Corollary 6 we could also use the algorithm from [34] which decides for given root sets whether there is an isomorphism mapping the root sets to each other.

We remark that our algorithm for the theorem does not necessarily compute a connected root set. In fact the number of connected root sets that yield distance decompositions of smallest width cannot be bounded by an fpt function, and they in particular cannot be enumerated in fpt time.

This can be seen as follows. We define the (k, p)-*path* as the graph obtained from P_k by replacing each edge with p internally vertex-disjoint paths of length 2 (see Figure 1). We call the original vertices of P_k *black* and the others *white*. The $(k+1, k)$-path is called k-path by Ding and Oporowski [12]. Among others, they use these graphs to characterize graphs of bounded strong tree width.

Theorem 11. *For every $k \geq 3$, there exists an infinite family of graphs such that each graph G in the family satisfies $\mathrm{rctdw}(G) = 2k - 1$ and has $((n-k)/(k-1))^{k-1}$ connected sets S for which $\mathrm{tdw}_S(G) = 2k - 1$, where n is the number of vertices of the graph.*

Proof. Let G be the (k, p)-path where $p \geq 4k - 2$. Let S be the vertices on a shortest path between the leftmost and rightmost black vertices of G. It is easy to see that there are p^{k-1} such sets S and $\mathrm{tdw}_S(G) = 2k - 1$. Since $p = (n - k)/(k - 1)$, it suffices now to show that $\mathrm{rctdw}(G) \geq 2k - 1$.

Assume that $\mathrm{tdw}_{S'}(G) < 2k - 1 \leq 4k - 2$ for some connected set $S' \subseteq V(G)$. Then, by Lemma 7 and the assumption $p \geq 4k - 2$, in any strong tree decomposition of width less than $2k - 1$, all black vertices have to be in the same bag. If they are in S', then $G[S']$ contains a path between the leftmost and rightmost black vertices. This implies $|S'| \geq 2k - 1$, a contradiction. Thus no black vertex can belong to S'. Now observe that every connected subgraph of G without black vertices has only one vertex. Therefore it is impossible to put all the black vertices into the same bag. □

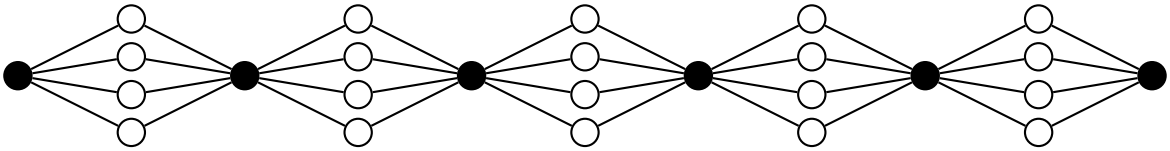

Figure 1: The $(6, 4)$-path.

The fact that the number of initial sets that achieve a decomposition of optimal width is not bounded by a fixed-parameter tractable functions prevents our isomorphism algorithm from being directly used to compute all optimal initial sets in fixed-parameter tractable time. In fact our isomorphism algorithm does not determine the root-connected tree distance width of the input graphs.

The same family of graphs mentioned above can be used to observe that the root-connected tree distance width of a graph cannot be bounded in terms of the rooted tree distance width (i.e., when only one vertex in the root is allowed).

Theorem 12. *For every $k \geq 2$, there is a graph G for which* $\mathrm{rctdw}(G) \leq 5$ *and* $\mathrm{rtdw}(G) \geq k$.

Proof. For any $p \geq 2k$, the $(3, p)$-path is an example of such a graph. □

Note that, in contrast to this, when we consider only path distance decompositions, the root-connected path distance width of a graph is bounded by a function of the rooted path distance width as shown in [27].

5 A reduction from strong tree width to maximum degree.

To define a reduction from isomorphism parameterized by strong tree width to isomorphism parameterized by maximum degree, we first reduce the problem to biconnected graphs relative to an equivalence relation and a suitable type of coloring compatible with the equivalence relation.

Let R be an equivalence relation on the vertices of a graph G. We define the *quotient graph* of G with respect to R as the graph whose vertex set consists of the equivalence classes and in which two equivalence classes E_1 and E_2 are adjacent if there exists vertices $v_1 \in E_1$ and $v_2 \in E_2$ such that v_1 and v_2 are adjacent. We define *the biconnected components of G relative to R* as the sets of vertices that comprise biconnected components of the quotient graph, i.e., pre-images of biconnected components under the projection to the quotient graph. A *tuple-coloring* is a map that assigns a color to every linear ordering of vertices in an equivalence class. When concerned with isomorphism of tuple-colored graphs equipped with an equivalence relation, we demand that isomorphisms preserve equivalence classes and the tuple-coloring, that is, an isomorphism must map an equivalence class to an equivalence class and colored ordered tuples to ordered tuples of the same color.

We define the *biconnected component tree* (also called block-cut tree) of a graph G with respect to R as the following bipartite graph: the vertices of the one partition class are those equivalence classes that form cut-vertices in the quotient graph. The vertices of the other partition class are the biconnected components of G relative to R. In the tree there is an edge between an equivalence class and a biconnected component if the corresponding cut-vertex is contained in the corresponding biconnected component in the quotient graph.

The quotient graph can be constructed in time linear in the number of edges of a graph. Since the biconnected components of the quotient graph can then also be computed in polynomial time, the biconnected components and the component tree with respect to R can be computed in polynomial time.

Lemma 13. *The isomorphism problem of graphs with an equivalence relation on the vertices Turing-reduces to isomorphism of tuple-colored biconnected components of the input graphs relative to the equivalence relation. The running time is bounded by a polynomial in n and $k!$, where n is the size of the input graphs and k is the size of the largest equivalence class.*

Proof. We describe an algorithm that has access to an oracle that can perform isomorphism queries for tuple-colored biconnected components of the input graphs relative to the equivalence relation. Let G_1 and G_2 be two input graphs each equipped with an equivalence relation. Since we can perform isomorphism tests for pre-images of components of the quotient graph separately, without loss of generality we can assume that the quotient graphs are connected. Moreover, we can assume that each graph has a distinguished biconnected component and restrict ourselves to finding isomorphisms that match the two components. Indeed, by fixing distinguishing an arbitrary biconnected component of G_1 and iterating over different possible choices for a distinguished component of G_2, we will find an isomorphism, if there exists one. We use the distinguished biconnected components as roots for the biconnected component trees.

Our technique from here on essentially performs dynamic programming on these rooted trees. More precisely, the algorithm proceeds as follows (details are given in Algorithm 3). If both input graphs are already biconnected, we use the oracle to determine isomorphism. If one input graph is biconnected while the other one is not, then the graphs are non-isomorphic. Thus, we suppose both input graphs are not biconnected. Consider the leaves of the biconnected component trees. Note that these cannot be cut-vertex classes and thus they are biconnected components. For each leaf L we proceed as follows. Let V_L be the unique cut-vertex class of L. We assign to every ordering σ of V_L a color that depends only on the isomorphism type of the graph induced by L in which σ has a special color. In other words, if two orderings of two cut vertex classes V_L and $V_{L'}$ obtain the same color, then there is an isomorphism of the corresponding biconnected components L and L' which maps the cut vertex classes to each other while observing their ordering. More formally we compute a function χ such that for leaves L and L' with orders σ and σ', respectively, we have $\chi(L,\sigma) = \chi(L',\sigma')$ if and only if said isomorphism exists. To determine this isomorphism type we use the oracle. Since we may have to check all pairs of leaves, this may require a number of oracle calls quadratic in the number of colored graphs whose isomorphism types are to be determined.

For every cut-vertex class adjacent to a leaf, we assign to every ordering of this cut-vertex class a new color that depends on its previous color (in case the ordering was already colored) and the multiset of colors computed for this ordering for all leaves. More precisely, the color depends on the multi-set $\{\{\chi(L,\sigma) \mid L \text{ a leaf and } \sigma \text{ is an ordering of } V_L\}\}$.

We then remove all vertices contained in leaves but not contained in cut-vertices from the input graphs.

To prevent the names of colors from becoming excessively long strings, we always rename newly arising colors by assigning previously unused integers to them.

After the modification, the graphs are isomorphic as tuple-colored graphs if and only if they were isomorphic before. By repeating the process at most a linear number of times the process will terminate with one of the graphs being biconnected in the last iteration.

Since in each iteration the modified graphs are isomorphic if and only if the original graphs are isomorphic, the algorithm correctly determines isomorphism. It remains to bound the running time. Since all steps between oracle calls can be performed in time polynomial in n and $k!$, it suffices to bound the number of oracle calls. The number of these is at most the number of pairs of biconnected components with distinguished cut-vertex classes equipped with a linear ordering. The number of biconnected components and cut-vertex classes is polynomial in n while the number of orderings of a cut-vertex class is at most $k!$. □

Algorithm 3 A reduction to biconnected components relative to an equivalence relation $\mathrm{Red}(G_1, G_2, R_1, R_2, B_1, B_2)$

Input: Two tuple-colored graphs G_1 and G_2 with equivalence relations R_1 and R_2, such that the quotient graphs are connected and for each a distinguished biconnected component B_1 and B_2, respectively. An Oracle Orc that determines isomorphism of biconnected graphs.

Output: **Yes** if there is an isomorphism from G_1 to G_2 that maps B_1 to B_2, **No** otherwise.

while both graphs G_1 and G_2 are not biconnected relative to R_1 and R_2, respectively **do**
  compute the biconnected component trees T_1 and T_2 with roots B_1 and B_2
  for all leaves L **do**
    compute the unique cut-vertex class V_L adjacent to L
  end for
  compute $M = \{(L, \sigma) \mid L \text{ is a leaf and } \sigma \text{ a linear order of } V_L\}$
  for all pairs of $\{((L_1, \sigma_1), (L_2, \sigma_2)) \mid (L_i, \sigma_i) \in M\}$ **do**
    use Orc to decide if there exists an isomorphism from L_1 to L_2 mapping V_{L_1} to V_{L_2} respecting the orders σ_i
  end for
  compute an isomorphism invariant coloring $\chi(L, \sigma)$ for every $(L, \sigma) \in M$
  assign to every ordering σ of a cut-vertex class V adjacent to a leaf a previously unused color that depends on the multi-set $\{\{\chi(L, \sigma) \mid (L, \sigma) \in M\}\}$
  remove all leafs L from both graphs
end while
if both graphs G_1 and G_2 are biconnected relative to R_1 and R_2 respectively **then**
  return $\mathrm{Orc}(G_1, G_2, B_1, B_2)$.
else
  return false
end if

From the theorem we obtain as corollary that to decide isomorphism of graphs in a hereditary graph class, i.e., a class closed under taking induced subgraphs, it suffices to be able to decide isomorphism of biconnected vertex-colored graphs.

Corollary 14. *The graph isomorphism problem of vertex-colored graphs in a hereditary graph class $\mathcal{C}$ polynomial-time Turing-reduces to the isomorphism problem of biconnected vertex-colored graphs in $\mathcal{C}$.*

Proof. The class of colored graphs is the same as the class of *tuple-colored* graphs where the equivalence relation is equality. Thus, the corollary follows from Lemma 13 since the algorithm described in its proof does not alter the equivalence relation of the input graphs. □

For general graph isomorphism, for every integer k, it is possible to reduce the isomorphism problem to isomorphism of k-connected graphs by simply adding universal vertices adjacent to all other vertices. However, for hereditary graph classes, under application of this technique or similar gadget constructions, the graphs may not necessarily remain within the class. In fact, if there is reduction to 3-connected graphs analogous to Corollary 14, then graph isomorphism would be polynomial-time solvable in general. This can be seen by considering the isomorphism-complete class of bipartite graphs in which in one bipartition class every vertex has degree at most 2. This class does not contain any 3-connected graphs that have components with more than 2 vertices.

We will now employ the relation $\equiv_{2k}$ defined in Section 4. However, this relation is not

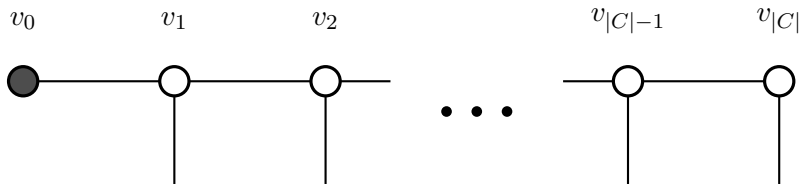

Figure 2: The gadget encoding the color of a linear ordering of a set C. The vertex v_0 is colored in the color that was associated with the linear ordering of C.

necessarily an equivalence relation. Let $\equiv_{2k}^{+}$ be the transitive closure of the relation $\equiv_{2k}$. It turns out that graphs that are biconnected relative to $\equiv_{2k}^{+}$ have bounded degree.

Lemma 15. *For $k \geq 2$, if a graph G with $\mathrm{stw}(G) \leq k$ is biconnected relative to $\equiv_{2k}^{+}$ then G has a maximum degree of at most $2k^2(k-1)+k-1$.*

Proof. The method of this proof is similar to the proof of Lemma 9. Suppose v is a vertex in G of degree $\deg(v)$ larger than $2k^2(k-1)+k-1$. Let $\mathcal{S}$ be a strong tree-decomposition of G of width at most k. By Lemma 7, vertices equivalent under $\equiv_{2k}$ must be in the same bag. Consequently vertices that are equivalent under $\equiv_{2k}^{+}$ must also be in the same bag. Since the degree of v is at least $2k$, the bag of v is adjacent to at least two bags which each contain a vertex u_1 and u_2 say. Moreover, the bag that contains v separates the vertices u_1 and u_2. Since the quotient graph is biconnected relative to $\equiv_{2k}^{+}$, from every vertex adjacent to but not equivalent to v under $\equiv_{2k}^{+}$ there is a path to u_1 and a path to u_2 that does not include any vertex equivalent to v under $\equiv_{2k}^{+}$. One of the paths must include a vertex from the bag of v.

Thus, for every vertex w adjacent to but not in the same bag as v, there is a path to a vertex contained in the bag of v that is not equivalent to v. By taking shortest paths that fulfill these conditions, we achieve that for neighbors of v in different bags, these paths are internally vertex-disjoint. At most $k-1$ vertices are contained in the same bag as v. Thus, since a bag contains at most k vertices, there are at least $(\deg(v)-(k-1))/k$ internally vertex disjoint paths from v to vertices which are in the same bag as v but not equivalent to v. Since $(|\deg(v)|-(k-1))/k > (k-1)2k$, by the pigeonhole principle, this shows the existence of a vertex in the same bag as v not equivalent to v $\equiv_{2k}^{+}$ from which there are $2k$ internally vertex-disjoint paths to v, yielding a contradiction. □

Lemma 16. *The isomorphism problem of tuple-colored graphs of degree at most d with an equivalence relation on the vertices with no equivalence class having more than k elements reduces to isomorphism of uncolored graphs of degree at most $\mathcal{O}(d+k!)$. The running time is polynomial in n and $k!$, where n is the size of the input graphs.*

Proof. We describe a reduction that satisfies the running time requirements to vertex-colored graphs. A standard reduction of attaching trees of equal height encoding vertex colors then reduces the problem further to uncolored graphs (see [28]).

Let C be an equivalence class. For every ordering of C, we attach a gadget to the vertices of C encoding the color of this ordering (see Figure 2). We add a path of newly added vertices $v_0, v_1, \ldots, v_{|C|}$. Vertex v_0 is colored with the color of the ordering of C. For $i \in \{1, \ldots, n\}$ an edge from v_i to the i-th vertex of C under the ordering is added. If this procedure is applied to two graphs, the resulting graphs are isomorphic if and only if the original graphs are. The maximum degree within the gadgets is 3, while the number of new edges added to a vertex originally in the graph is at most $k!$, where k is the size of the largest equivalence class. □

For the lemma, it is essential that equivalence classes do not intersect. By coloring partially overlapping sets, it is possible to encode hypergraphs, and thus to encode graphs, even with sets of size at most 2.Consequently a reduction for this case would yield a polynomial time algorithm for graph isomorphism.The construction in the previous proof yields a constructive method for representing a given permutation group, which must be possible by Frucht's Theorem [16]. In the meantime, extensive research on representablility of groups as automorphism groups of graphs has been conducted (see [9]).

Together, the lemmas of this section can be used to assemble a reduction from strong tree width to maximum degree.

Theorem 17. *There is an fpt Turing-reduction from isomorphism parameterized by strong tree width to isomorphism parameterized by maximum degree.*

Proof. We first compute $\equiv_{2k}$ and the transitive closure $\equiv^{+}_{2k}$ which can be done in polynomial time. If the largest equivalence class has size greater than k we reject the input as infeasible containing a graph of strong tree width larger than k. We then reduce via Lemma 13 to biconnected graphs relative to $\equiv^{+}_{2k}$. By Lemma 15 the biconnected components have bounded degree. By Lemma 16 we can then reduce the isomorphism problem of the tuple-colored biconnected components to isomorphism of graphs of bounded degree. □

6 Applications to fpt isomorphism results

In this section we combine the two results from the previous sections to obtain further fpt isomorphism algorithms.

Theorem 18. *Graph isomorphism parameterized by connected strong tree width can be solved in fpt time.*

Proof. By Theorem 17 the problem reduces to isomorphism of the tuple-colored biconnected components. Let G be a graph of connected strong tree width at most k of bounded degree. To apply Corollary 6 we first describe a set of possible bags capturing a strong tree decomposition of G computable in fpt time and having an fpt size bound. For this consider the family $\mathcal{V}(G)$ of sets of size at most k that project to a connected subgraph in the quotient graph relative to $\equiv_{2k}$. Since the degree of G is bounded, the number of such sets is bounded by an fpt number. The theorem now follows from Corollary 6. □

We have shown fixed-parameter tractability for isomorphism with respect to the parameters connected strong tree width and root-connected tree distance width. In turns out that these parameters are unrelated, i.e., that neither of these parameters can be bounded by a function of the other.

Theorem 19. *For every $k \geq 2$, there is a graph G for which* $\mathrm{rctdw}(G) \leq 2$ *and* $\mathrm{cstw}(G) \geq k$.

Proof. The cycle C_{2k} on $2k$ vertices has these properties. □

The (k,p)-*comb* is the graph obtained from the path P_k by attaching a copy of $K_{2,p}$ to each vertex of P_k as depicted in Figure 3. We call the vertices of degree 2 *white* and the others *black*. It is easy to see that for every k and p, the (k,p)-comb has connected strong tree width at most 3 (see Figure 4).

Theorem 20. *For every $k \geq 3$, there is a graph G for which* $\mathrm{cstw}(G) \leq 3$ *and* $\mathrm{rctdw}(G) \geq k$.

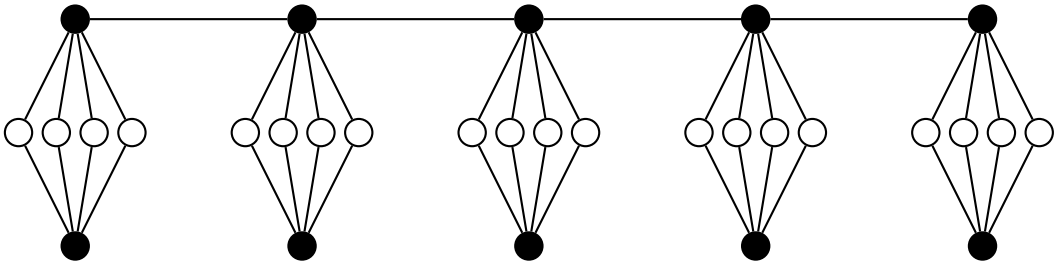

Figure 3: The $(5, 4)$-comb.

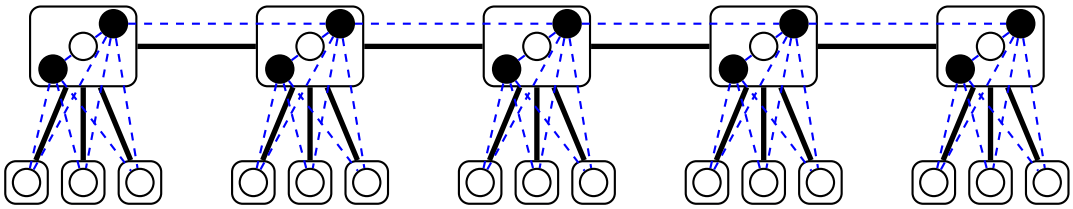

Figure 4: A connected strong tree decomposition of the $(5, 4)$-comb of width 3.

Proof. Let G be the (k, p)-comb where $p \geq 2k$. Suppose that $\mathrm{tdw}_S(G) < k$ for some connected set $S \subseteq V(G)$. Since $|S| < k$, there is a copy of $K_{2,p}$ in G that does not intersect S. Therefore, the two black vertices in that copy cannot be in the same bag. This contradicts Lemma 7. □

Lemma 21. *Let G be a connected graph with strong tree width at most k and for which the longest geodesic cycle has length at most c. There exists a strong tree decomposition of width at most k with the following property: If S is a proper subset of a bag B then there is a vertex v in $B \setminus S$ that is of distance at most c from S.*

Proof. Let G be a graph with strong tree width at most k. Consider a strong tree decomposition of width at most k for which the bag partition is finest. This implies that the components of the forest obtained by removing a bag induce connected graphs. We show that such a decomposition has the desired property. Thus let S be a proper subset of a bag B. If S has a neighbor in B this claim is obvious, so we suppose otherwise. If B were a leaf, then B would be connected, so we also suppose B is not a leaf. In this case, let B' be a neighbor of B whose subtree connects a vertex s from S with some vertex v in B outside of S. In the subtree of B in the graph in which B' has been deleted, there is a path from s to v. This implies there is a cycle containing a vertex in S and a vertex of B not in S. Let C be a shortest such cycle. We claim that this cycle is geodesic. If C were not geodesic, there would exist a shortcut connecting two vertices c_1 and c_2 of C via a path P shorter than the distance from c_1 to c_2 on the cycle. Since C was chosen of minimal length, this implies that the path P contains a vertex x of B. Together with two segments of C the path P forms two cycles shorter than C. Since P contains a vertex from B, one of the cycles contains both a vertex from S and $B \setminus S$. This contradicts minimality of C. □

As further examples of applications of our technique we can obtain fixed-parameter tractability results of other parameters as follows. A geodesic cycle in a graph G is a cycle C such that the distance between every two vertices in C is the same as the distance in G. The chordality of a graph is the length of the longest induced cycles.

Theorem 22. *Graph isomorphism parameterized by the maximum of the length of a geodesic cycle and strong tree width can be solved in fpt time.*

Proof. For a biconnected component C of G relative to $\equiv_{2k}$ we let $\mathcal{V}_C$ be the family of vertex sets B that have the property that for every subset $S \subset B$ there are vertices $s \in S$ and $v \in B \setminus S$ such that s and v are equivalent or s and v are of distance at most c where c is the length of the longest geodesic cycle of G. By the previous lemma this family contains all bags of a strong tree decomposition of C and by Lemma 15 this set is of fpt size. Moreover it can be computed in fpt time. We conclude the theorem with the same arguments already used in the proof of Theorem 18. □

Corollary 23. *Graph isomorphism parameterized by the maximum of the chordality and degree can be solved in fpt time.*

Proof. The corollary follows from the previous theorem by noting that bounded chordality implies a bound on the length of geodesic cycles. Moreover if the chordality and the degree is bounded, then the graph has bounded tree width [5], [24] and consequently it has bounded strong tree width [11], [33]. □

Further applications of the theorems can be obtained by considering the set of potential maximal cliques. A *potential maximal clique* of a graph G is a set of vertices that is a bag in some minimal tree-decomposition of G. The set of potential maximal cliques can be computed in polynomial time in the size of the set itself [6]. Moreover, this set is isomorphism invariant and the subset of potential maximal cliques of size k captures a tree decomposition of G of minimal width. We can thus apply Corollary 6 to the potential maximal cliques.

Theorem 24. *If for a parameterized graph class there is an fpt bound on the number of potential maximal bags then isomorphism is fixed parameter tractable in the maximum of the parameter and the tree width.*

Proof. By Corollary 6, it suffices to compute for all graphs an invariant fpt size family of sets capturing a tree decomposition.

The set of potential maximal cliques captures a tree decomposition, is isomorphism invariant, can be computed in time polynomial in its size [6] and is bounded by an fpt function by assumption. It thus satisfies the requirements. □

Equivalently, in the theorem it suffices to have a bound on the number of minimal s-t-separators, since by a theorem of Bouchitté and Todinca [6] the number of potential maximal cliques is polynomially bounded the number of minimal s-t-separators.

Using results from [15] we can apply the theorem to various graph classes. Indeed, it is known that the number of minimal s-t-separators is polynomially bounded for weakly chordal, polygonal circle, circular-arc and d-trapezoid graphs (see [15, Section 5]). So for all these classes, isomorphism is fixed parameter tractable when parameterized by tree width. Moreover, for weakly chordal graphs and circular-arc graphs the tree width of H-minor free graphs is bounded by a function of the number of vertices in H (see also [15, Section 5]), so isomorphism of H-minor free weakly chordal graphs and H-minor free circular-arc graphs is fixed parameter tractable when parameterized by the size of H.

7 Conclusion

In this paper we show that, in order to perform isomorphism tests, it suffices to compute an invariant set of potential bags that is comprehensive enough to express a tree decomposition. Indeed, by applying the restricted Weisfeiler-Lehman, we do not need to worry about how to perform the isomorphism test, nor how to compute a decomposition.

For various other results this means that their isomorphism testing part can be replaced by the general theorem and only the part analyzing the graph class remains. For example in [26] and [22] it is shown that for chordal graphs of bounded clique number a tree model can be computed in cubic time. This tree model is unique up to the ordering of children and gives rise to an invariant family of sets of vertices capturing a tree-decomposition.

In [7] it is shown that for graphs of tree-depth at most k the number of vertices that can be chosen as the root in a tree-depth decomposition is bounded by a function of k, recursively applying this gives rise to invariant family of sets of vertices capturing a tree-decomposition.

Furthermore we demonstrated how the theorems can be used in conjunction with the set of potential maximal cliques. The advantage here is that it is known in general that this set can be efficiently computed.

However, our technique can also be applied to tree decompositions, where the long question whether graph isomorphism is fixed parameter tractable when parameterized by tree width remains.

On the other hand, our parameterized reduction to bounded degree is only valid for strong tree width, and whether such a reduction exists for tree width remains open. Finally, as mentioned in the introduction, no non-tractability results are known in this context, and parameterized reduction could be a method to establish a hardness criteria.

References

[1] V. Arvind, B. Das, J. Köbler, and S. Toda. Colored hypergraph isomorphism is fixed parameter tractable. In *FSTTCS*, pages 327–337, 2010.

[2] C. Berkholz, P. Bonsma, and M. Grohe. Tight lower and upper bounds for the complexity of canonical colour refinement. In *ESA*, pages 145–156, 2013.

[3] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.

[4] H. L. Bodlaender. Necessary edges in k-chordalisations of graphs. *J. Comb. Optim.*, 7(3):283–290, 2003.

[5] H. L. Bodlaender and D. M. Thilikos. Treewidth for graphs with small chordality. *Discrete Applied Mathematics*, 79(1–3):45–61, 1997.

[6] V. Bouchitté and I. Todinca. Listing all potential maximal cliques of a graph. *Theor. Comput. Sci.*, 276(1-2):17–32, 2002.

[7] A. Bouland, A. Dawar, and E. Kopczynski. On tractable parameterizations of graph isomorphism. In *IPEC*, pages 218–230, 2012.

[8] J.-Y. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.

[9] P. J. Cameron. Automorphism groups of graphs. In *Selected topics in graph theory, 2*, pages 89–127. Academic Press, London, 1983.

[10] S. Datta, N. Limaye, P. Nimbhorkar, T. Thierauf, and F. Wagner. Planar graph isomorphism is in log-space. In *IEEE Conference on Computational Complexity*, pages 203–214, 2009.

[11] G. Ding and B. Oporowski. Some results on tree decomposition of graphs. *J. Graph Theory*, 20(4):481–499, 1995.

[12] G. Ding and B. Oporowski. On tree-partitions of graphs. *Discrete Math.*, 149(1–3):45–58, 1996.

[13] S. Evdokimov and I. N. Ponomarenko. Isomorphism of coloured graphs with slowly increasing multiplicity of jordan blocks. *Combinatorica*, 19(3):321–333, 1999.

[14] J. Flum and M. Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer, London, UK, March 2006.

[15] F. V. Fomin, I. Todinca, and Y. Villanger. Large induced subgraphs via triangulations and cmso. In *SODA*, pages 582–583, 2014.

[16] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compos. Math.*, 6:239–250, 1938.

[17] F. Fuhlbrück. Fixed-parameter tractability of the graph isomorphism and canonization problems. Diploma thesis, Humboldt-Universität zu Berlin, 2013.

[18] M. L. Furst, J. E. Hopcroft, and E. M. Luks. Polynomial-time algorithms for permutation groups. In *FOCS*, pages 36–41, 1980.

[19] M. Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. In *LICS*, pages 179–188, 2010.

[20] M. Grohe and J. Mariño. Definability and descriptive complexity on databases of bounded tree-width. In *ICDT*, pages 70–82, 1999.

[21] N. Immerman and E. Lander. *Complexity Theory Retrospective*, chapter Describing Graphs: A First-Order Approach to Graph Canonization. Springer-Verlag, 1990.

[22] M. Klawe, D. Corneil, and A. Proskurowski. Isomorphism testing in hookup classes. *SIAM Journal on Algebraic Discrete Methods*, 3(2):260–274, 1982.

[23] J. Köbler, U. Schöning, and J. Torán. *The graph isomorphism problem: its structural complexity*. Birkhäuser Verlag, Basel, Switzerland, Switzerland, 1993.

[24] A. Kosowski, B. Li, N. Nisse, and K. Suchan. k-chordal graphs: From cops and robber to compact routing via treewidth. In *ICALP (2)*, pages 610–622, 2012.

[25] S. Kratsch and P. Schweitzer. Isomorphism for graphs of bounded feedback vertex set number. In *SWAT*, pages 81–92, 2010.

[26] T. Nagoya. Counting graph isomorphisms among chordal graphs with restricted clique number. In *ISAAC*, pages 136–147, 2001.

[27] Y. Otachi. Isomorphism for graphs of bounded connected-path-distance-width. In *ISAAC*, pages 455–464, 2012.

[28] P. Schweitzer. *Problems of unknown complexity: graph isomorphism and Ramsey theoretic numbers*. Phd thesis, Universität des Saarlandes, Saarbrücken, Germany, July 2009.

[29] P. Schweitzer. Isomorphism of (mis)labeled graphs. In *ESA*, pages 370–381, 2011.

[30] D. Seese. Tree-partite graphs and the complexity of algorithms. In *Fundamentals of Computation Theory*, FCT '85, pages 412–421, 1985.

[31] S. Toda. *Gurafu Doukeisei Hantei Mondai (The Graph Isomorphism Decision Problem)*. Nihon University, Tokyo, Japan, 2001. in Japanese.

[32] S. Toda. Computing automorphism groups of chordal graphs whose simplicial components are of small size. *IEICE Transactions*, 89-D(8):2388–2401, 2006.

[33] D. R. Wood. On tree-partition-width. *Eur. J. Comb.*, 30(5):1245–1253, 2009.

[34] K. Yamazaki, H. L. Bodlaender, B. de Fluiter, and D. M. Thilikos. Isomorphism for graphs of bounded distance width. *Algorithmica*, 24(2):105–127, 1999.